intellibooks

# WHAT IS AI GOVERNANCE?

The framework every org needs before scaling AI

## 1. THE EVOLUTION

**Stage 1 (2015)**
Ad-hoc AI
Uncontrolled experimentation

**Stage 2 (2020)**
Policy-Driven AI
Basic rules & approvals

**Stage 3 (2025)**
Scaled, Governed AI
Governance embedded across all AI systems

## 2. THE 4 CORE LAYERS

**RISK CLASSIFICATION**
- Classify AI use cases by impact (low / medium / high)
- Define allowed, restricted, and prohibited

**MODEL ACCOUNTABILITY**
- Assign clear ownership for every deployed model
- Track versions, training data, and intended use

**MONITORING & AUDITABILITY**
- Log all model decisions with traceability
- Detect drift, bias, and performance issues in real time

**HUMAN OVERSIGHT**
- Define when human intervention is required
- Establish escalation paths for failures and edge cases

## 3. WITHOUT VS WITH GOVERNANCE

| ✕ Without AI Governance | ✓ With AI Governance |
|---|---|
| No visibility into AI decisions | Complete audit trail of decisions |
| Unowned, unmanaged models | Defined ownership and accountability |
| Late detection of compliance risks | Proactive risk management |
| Inconsistent AI behavior | Standardized processes across teams |
| Shadow AI usage | Controlled and approved AI ecosystem |
| High reputational risk | Prepared incident response strategies |

## 4. TOOL LANDSCAPE

| Tool | Approach | Coverage | Strength |
|---|---|---|---|
| Microsoft Purview | Data + AI governance | Organization level | Data + AI governance |
| IBM watsonx governance | Bias & drift monitoring | Model-level | Bias & drift monitoring |
| Aporia | Real time guardrails | Inference level | Real time guardrails |
| Fiddler AI | Explainability + monitoring | Model level | Explainability + monitoring |
| Credo AI | Policy as code | Portfolio level | Policy as code |

## 5. THE MODERN GOVERNANCE STACK

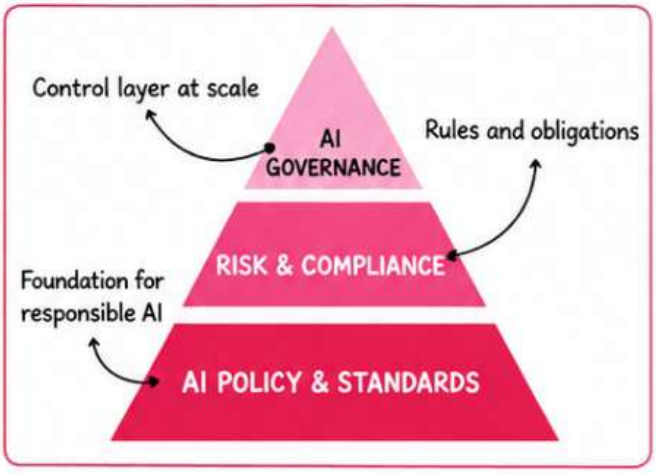


**AI GOVERNANCE:** Often ignored until issues arise

**RISK & COMPLIANCE :** Where many initiatives slow down

**AI POLICY & STANDARDS:** The essential starting point

## 6. QUICK WINS

- Build an AI use case registry (track all AI usage)
- Assign an owner to every production model
- Create a one-page AI Acceptable Use Policy
- Add human review for high-risk AI outputs
- Audit for shadow AI tools in use
- Choose a framework (NIST or EU AI Act) and identify gaps